Method for access to a medium by a multi-channel device

The invention relates to methods for access to a medium by a multi-channel device. The medium comprises a transmission system having at least two channels on which a message to be transmitted comprises at least a preamble and a header plus a succeeding data field. The data field may contain either user data, which is termed "useful data", or

5   control information, for coordinating the access to the medium for example. A method for decentralized medium access and a method for centralized medium access are claimed.

The frequency band of a transmission system is often divided into sub-bands, which are termed channels, on which either a single communications link or else a complete cell of a system operates. The latter is the case with for example 802.11a/e and HiperLAN/2

10   wireless local area networks (WLANs). With regard to 802.11 a/e WLANs, the basic specification, namely ANSI/IEEE Std 802.11, 1999 edition, and "IEEE Std 801.11a-1999: High-speed Physical Layer in the 5 GHz Band, Supplement to Standard IEEE 802.11", IEEE New York, Sep. 1999 and "IEEE Std 802.11e/D4.2: Medium Access Control (MAC) Enhancements for Quality of Service (QoS), Draft Supplement to Standard IEEE 802.11",

15   IEEE New York, Feb 2003, are hereby incorporated by reference.

The bandwidth of the channel sets a limit to the maximum data rate that can be obtained between two stations in a cell, or in other words to the capacity of a cell. One

20   possible way of increasing the capacity of a transmission system is to enlarge the bandwidth of a communications channel. With a preset channel definition, this can be achieved by grouping two or more channels to obtain one channel of greater width. This approach is known in theory and has been implemented in certain WLAN systems, in which a high data rate is obtained or a what is called "turbo" mode of handling.

25   There are a certain number of stations in a transmission system, and of these two or more at a time make a temporary connection with one another. The grouping of channels can only be performed if all the stations involved operate in the mode where the data rate is high. For channel grouping, there has to be one standard for all the stations or terminals. However, mobile telecommunications is the very area in which there are a variety

of manufacturers of terminals and the terminals employ channel grouping that is not standardized. Because channel grouping, as a performance feature, entails additional expenditure in development and manufacture, it may also be the case that, because of the cost, a manufacturer of terminals offers on the one hand terminals employing channel

5   grouping and on the other hand terminals that do not have this performance feature. Also, older terminals that were developed and sold before channel grouping was introduced in transmission systems are incapable of operating in the mode having the high data transmission rate.

Standard IEEE 802.11e combines both a decentralized and a centralized

10   scheme for medium access. An algorithm that can be used for this will be explained by taking the 802.11 system as an example. The basic 802.11 MAC protocol is the distributed coordination function (DCF), that operates as a listen-before-talk scheme. The distributed coordination function is based on carrier multiple sense access (CMSA). Having found that no other transmissions are underway on the wireless medium, the stations emit MAC service

15   data units (MSDUs) of an arbitrary length of up to 2304 bytes. If, however, two stations find that a channel is free at the same time, a collision occurs when the data is sent out over the radio-transmitting medium, namely air, that they both use. The 802.11 MAC protocol defines a mechanism for collision avoidance (CA) to reduce the probability of collisions of this kind. It is part of the collision avoidance mechanism that, before it starts transmitting, a station

20   performs a waiting or backoff process. The station continues listening out on the channel for an additional, random period of time after it finds that the channel is free. Only if the channel remains free for this additional random length of time is the station permitted to initiate a transmission. This random waiting time is composed of a constant portion, the what is called DCF interframe space (DIFS), which is 34 µs long in the case of the 802.11a MAC protocol,

25   and a random portion whose length is between zero and a maximum time. The DIFS space is thus the minimum possible waiting time for the stations. The length of the random portion of the waiting time is obtained as a multiple of the length of a time slot (slot time), which length is 9 µs in the 802.11a MAC protocol. Each station draws a random value, for the number of slot times to be waited, which it stores in the what is called contention window (CW). On the

30   expiry of each period of 9 µs, the value of the CW is decremented by 1.

Each time a data frame is successfully received, the receiving station immediately sends out an acknowledgement frame (ACK). The size of the contention window is enlarged if a transmission fails, which means that a data frame sent out has not been acknowledged. After each failed attempt at a transmission, a new medium access is

effected after a fresh waiting time, with the fresh waiting time being selected to be twice as long as the current contention window. This reduces the likelihood of a collision in the event that a plurality of stations are trying to gain access to the channel. Those stations that deferred channel access during the time when the channel was busy do not select a new

5   random waiting time but continue the countdown of the time for the deferred medium access on finding that the channel is idle again. In this way, stations that deferred channel access due to their longer random waiting time as compared with other stations are given a higher priority when they resume their efforts to start a transmission. After each successful transmission, the transmitting station performs a new random waiting process (backoff) even

10   if it does not have a further MSDU to send at the time. This is what is called the "post backoff", because this waiting process takes place after rather than before a transmission.

There is a situation under the 802.11 MAC protocol in which a station does not have to perform a waiting process of random duration (a backoff) before it can start transmitting data. This situation arises if an MSDU from a higher layer arrives at a station

15   and the post-backoff for the last transmission has already been completed, or in other words, there is no queue and, in addition, the channel has been idle for a minimum DIFS time. All subsequent MSDUs that arrive after this MSDU will be transmitted after a random waiting time until there is again no queue.

To limit the probability of long frames colliding and being transmitted more

20   than once, data frames are also fragmented. A long MSDU can be divided by fragmentation into a plurality of small data frames, i.e. fragments, that can be transmitted sequentially as data frames to be acknowledged individually. The advantage of fragmentation is that, if a transmission fails, this failure can be detected at an earlier point in time and hence less data has to be re-transmitted.

25   In systems using CSMA, there is a problem with hidden stations. This is a problem inherent in the CSMA system and to alleviate it the system defines a request-to-send/clear-to-send (RTS/CTS) mechanism that can be used as an option. Before data frames are transmitted, it is possible for a system to send a short RTS frame, which is followed by a CTS transmission from the receiving station. The RTS and CTS frames contain information

30   on the length of the transmission time of the next data frame, i.e. the first fragment, and of the corresponding ACK response. What is achieved in this way is that other stations near the transmitting station, and hidden stations near the receiving station, do not start a transmission, because they set a counter, the what is called Network Allocation Vector (NAV). The RTS/CTS mechanism helps to protect long data frames against hidden stations. With

fragmentation, a large number of ACKs are transmitted, whereas with RTS/CTS the MSDU can be transmitted efficiently in a single data frame. Between each successive pair of frames in the sequence RTS frame, CTS frame, data frame and ACK frame, there is a short interframe space (SIFS), which is 16 µs long under 802.11a.

5   Fig. 1, which relates to the prior art, is a diagram showing an example of a distributed coordination function (DCF). A short interframe space (SIFS) is shorter than a DCF interframe space (DIFS), as a result of which CTS responses and acknowledgement frames (ACKs) always have the highest priority for access to the wireless medium. In the latest version of the MAC protocol, the 802.11e protocol, an Enhanced Distribution

10   Coordination Function (EDCF) has been introduced, which still operates in the same way but, in addition, supports different types of traffic, such as, for example, access priorities. In the time-based diagram for six stations shown in Fig. 1, although station 6 cannot detect the RTS frame of the station 2 that is transmitting, it can detect the CTS frame of station 1.

          Another known function, the Hybrid Coordination Function (HCF) extends the

15   rules for (E)DCF access. A crucial performance feature of 802.11e MAC is the Transmission Opportunity (TXOP). A TXOP is defined as the interval between the point at which a station receives the right to initiate a transmission, defined by the starting time, and a maximum duration. TXOPs are allocated by way of contentions (EDCF-TXOP) or are granted by HCF (polled TXOP). Only one station in the cell, which is called the hybrid coordinator (HC), can

20   give other stations permission to transmit, i.e. can grant a TXOP. The duration of a polled TXOP is specified by the time field within the allocating frame. The hybrid coordinator is able to allocate TXOPs to itself, to enable MSDU transmissions to be initiated, at any time, but only on detecting that the channel is idle for a time equal to a PIFS (Point Coordinator Interframe Space), which time is shorter than the length of a DIFS.

25   Defined as part of the 802.11e protocol is an additional random access protocol that enables collisions to be reduced. What is called Controlled Contention gives the hybrid coordinator an opportunity of learning what stations need to be queried at what times with regard to their wishes to transmit. The controlled contention mechanism allows stations to request to be allocated polled TXOPs by sending a source query, without interfering with

30   other (E)DCF traffic.

          It is an object of the invention to specify methods for medium access by a multi-channel device on a transmission system having at least two channels, which methods

give terminals not having the performance feature of channel grouping the opportunity of transmitting and receiving on said transmission systems, i.e. methods that make it possible for single-channel and multi-channel devices to coexist. Methods are to be specified for a centralized mechanism and a decentralized mechanism.

5.   The object is achieved in accordance with the invention by a method for access to a medium by a multi-channel device, which medium comprises a transmission system having at least two channels, on which a message to be transmitted comprises at least a preamble, a header and a succeeding control or data section of a frame, and the preamble and header of the message are repeated on all the channels. As a result of the control

10   information in a frame being repeated on every channel, even a single-channel device is able to pick up the preambles and headers and perform a standards-compliant waiting process, possibly with a random backoff time. If the PHY and the MAC headers were not repeated on every frequency channel, single-channel devices would not recognize the transmission and would interpret it is as interference that blocks their channel for an indeterminate length of

15 . time.

              In one embodiment of the invention, the preamble and header are repeated in parallel on all the channels. The parallel transmission of these parts of the message cannot begin until all the channels are free. Simulations have shown that the loss in data transmission rate that occurs as a result of the parallel transmission is relatively slight.

20   In a further embodiment, the succeeding control or data section of a frame is taken from the group request-to-send (RTS), clear-to-send (CTS), acknowledgement (ACK) or data (DATA).

              In a special embodiment of the invention, the multi-channel device operates to standard IEEE 802.11, i.e. 802.11e or 802.11n having a medium access control (MAC)

25   protocol, and it is not only the preambles and headers that are repeated on all the channels but also at least some of the items of information belonging to the MAC protocol.

              If the medium access takes place under standard IEEE 802.11, i.e. 802.11e, the RTS, CTS and ACK control frames are transmitted on all the channels, which means that the single-channel devices set their network allocation vectors (NAVs) on the basis of the

30   information in the RTS/CTS data packets. The setting of the NAV causes a counter to start, and no attempt is made to access the radio transmission medium until a target count has been reached.

              The object is also achieved in accordance with the invention by a method for access to a medium by a multi-channel device, which medium comprises a transmission

system having at least two channels that the multi-channel device intends to call upon for transmission, which method has the following steps

- scanning by the multi-channel device of all the channels to be called upon for transmission,

5   - finding that a single one of these channels is idle or that a back-off by the device itself is underway on this channel,

- blocking of this channel to other devices by the multi-channel device,

- further scanning of the other channels to be called upon and blocking or reserving thereof on finding that the channel concerned is idle or that a backoff is underway

10   thereon.

What is finally achieved as a result of the successive blocking of individual channels is a state where all the channels to be called upon for transmission are idle. Other devices that scan a blocked channel recognize that the channel is not idle and therefore do not themselves begin a transmission. The transmission of the message can then be performed

15   with grouped channels and hence at a high data rate.

In one embodiment, the blocking of a channel that is idle is performed by the multi-channel device and the receiving device, each of which emits a reserving message.

In one embodiment, the reserving message is implemented in the form of RTS and CTS frames that are transmitted by the following steps

20   - transmission of an RTS frame on the free channel by the multi-channel device, so that devices in the area surrounding the multi-channel device that is transmitting will set their NAVs,

- transmission of a CTS frame on the free channel by the receiving device, so that stations in the area surrounding the receiving station will set their NAVs.

25   The multi-channel device is able to carry out its transmission with channel grouping in this case, on all the channels that it has itself previously blocked.

In one embodiment, the blocking of a channel that is idle for a predetermined period of time is performed by starting the transmission by the multi-channel device on the single channel, doing so alternatively with or without an RTS-CTS mechanism. What the

30   RTS-CTS mechanism entails is

- optional transmission of an RTS frame on the free channel by the multi-channel station,

- optional transmission of a CTS frame by the receiving station and

- transmission of a data frame on the free channel by the multi-channel station (and hence blocking of the channel).

The object is also achieved in accordance with the invention by a method for access to a medium by a multi-channel device, which medium comprises a transmission system having at least two channels that the multi-channel device intends to call upon for transmission, in which method a third device (independent of the transmitter and receiver) reserves or blocks the channels in the channel group for the multi-channel device that wishes to transmit. This third device has assumed responsibility on the network for emitting the synchronizing beacon.

In one embodiment, the third device is responsible for coordinating medium access to a plurality of channels.

In one embodiment, in the event of individual channels in the channel group not becoming free simultaneously, the third device causes, alternatively

(a)   one channel or individual channels to be blocked until such time as all the channels in the channel group have become free, or

(b)   a channel that has become free to be assigned immediately to the multi-channel device that wishes to transmit.

In one embodiment, the medium access is performed under standard IEEE 802.11, i.e. 802.11e or 802.11n, and said third device is generally the so-called hybrid coordinator or point coordinator. The centralized reservation of the channels is performed by the third device, which is responsible for emitting the beacon on all the channels in the channel group and at the same time keeps free or reserves all the channels in the channel group for the multi-channel station that wishes to transmit. If said third device or station is responsible for coordinating medium access to all the channels, it can make the reservations for the devices on all the channels in such a way that phases able to be used simultaneously are provided for the multi-channel devices on all the channels in their channel groups. Should the individual channels in the channel group nevertheless not become free simultaneously, there are two different ways in which said third station is able to secure the channel for the multi-channel station until all the channels in the channel group become free.

In a preferred embodiment of the method, the medium access is performed under standard IEEE 802.11, i.e. 802.11e or 802.11n, and said third device is the hybrid coordinator or point coordinator.

In a further embodiment of the method, the point coordinator or hybrid coordinator transmits what is called beacons in parallel on all the channels. The transmission may if required be performed simultaneously or in synchronized fashion.

The method in accordance with the invention can be employed on a

5   transmission system using the Standard Universal Mobile Telecommunication System (UMTS).

These and other aspects of the invention are apparent from and will be elucidated with reference to the embodiments described hereinafter.

10

In the drawings:

Fig. 2 is a time-based diagram of a prior art transmission.

Fig. 3 is a time-based diagram of a first variant of the method in accordance with the invention.

15   Fig. 4 is a time-based diagram of a second variant of the method in accordance with the invention.

Fig. 2 is a time-based diagram of a prior art transmission. The transmission

20   always takes place on only one channel even when a channel is idle, as channel 3 (ch 3) is in the present case.

Fig. 3 is a time-based diagram of a first variant of the method in accordance with the invention. Possible examples of messages are shown on individual ones of six channels. Two stations usually communicate on a channel by alternately transmitting an RTS

25   frame, a CTS frame, a DATA frame and an ACK frame. At time T_1, a multi-channel device intends to start transmitting data using channel grouping. In the present example, scanning of the six channels shows that channels 1, 2 and 5 are idle, whereupon, following a preamble and header PR, an RTS frame is first transmitted and a CTS frame is then received. On reception of the three CTS frames, which are parallel in this case, channels 1, 2 and 5 are

30   blocked, which means that any other stations, regardless of whether they are single-channel or multi-channel devices, have set their NAVs. At time T_2, the multi-channel device detects that channel 4 is idle, whereupon the RTS-CTS procedure is performed and this channel is thus blocked. The same thing happens with channel 3 at time T_3. Once five out of six channels have been blocked, the transmission of data starts at time T_4, when the final channel

too, namely channel 6, is recognized to be idle. The preamble and header PR having been transmitted on the six individual channels, the transmission of the data DATA takes place using channel grouping, by which means a higher data rate is attained. On completion of the transmission of the data, the transmitting multi-channel device receives on each individual

5   channel, following a preamble and header PR, an acknowledgement frame ACK. An advantage of this variant is that the transmission of data starts at a fixed, high data rate right at the outset and does so simultaneously on all the channels.

Fig. 4 is a time-based diagram of a second variant of the method in accordance with the invention. In this example, the messages and the busy/idle states on the channels

10   before the scanning by the multi-channel device begins are assumed to be the same as in Fig. 3. However, on channels that are in the idle state, the transmission of data is begun for the purpose of blocking them, which happens at time T_4 in this case. On other channels, transmission begins (possibly after a pause) as soon as it is detected that a channel is idle, which is detected to be the case on channel 4 at time T_5 and on channel 3 at time T_6 in the

15   present case. At time T_7, the final channel, namely channel 6, is brought in as well. On this channel, the message first begins with a preamble and header. Following this, channel grouping then takes place for the DATA frame using all six channels. At the beginning T_4 of the transmission of the data, three individual channels are used in parallel and only at time T_7 is the bandwidth increased by channel grouping.

20   It is true of all three diagrams that a space is shown between the end of a control frame and the start of a preamble and header, which space is intended to show that there may be a brief pause here.